

UPPER BOUNDS ON PAIRS OF DOT PRODUCTS

DANIEL BARKER AND STEVEN SENGER

ABSTRACT. Given a large finite point set, $P \subset \mathbb{R}^2$, we obtain upper bounds on the number of triples of points that determine a given pair of dot products. That is, for any pair of positive real numbers, (α, β) , we bound the size of the set

$$\{(p, q, r) \in P \times P \times P : p \cdot q = \alpha, p \cdot r = \beta\}.$$

1. INTRODUCTION

Many elementary problems in geometric combinatorics ask how often a particular type of point configuration can occur in subsets of some ambient space. One of the most famous is the Erdős single distance problem, which asks how often any fixed distance can occur in a large finite set of points in the plane. The conjecture is that for a set of n points, no distance can occur more than $Cn^{1+\epsilon}$ times, for some constant C , independent of n , and any $\epsilon > 0$. The best known estimate of this is due to Spencer, Szemerèdi, and Trotter, in [11], who have shown that $Cn^{\frac{4}{3}}$ is an upper bound. A closely related problem, the Erdős distinct distances problem, asks for a lower bound on the number of distinct distances determined by point pairs from a large finite point set. This was resolved in the plane by Guth and Katz, in [5]. Analogous questions have been studied for dot products; see [8], [12], and [4].

Here, we consider triples of points which determine a pair of dot products in large finite point sets. In the settings of vector spaces over various finite rings, there has been activity on the special case of zero dot products by the second listed author, and Iosevich [9], as well as Pham and Vinh, in [14]. Information about the dot products determined by a point set finds applications in varied areas such as coding theory, [1], graph theory, [2], and frame theory, [3].

We now fix some notation. In what follows, if two quantities, X and Y , vary with respect to some natural number parameter, n , then we write $X \lesssim Y$ if there exist constants, C and N , both independent of n , such that for all $n > N$, we have $X \leq CY$. If $X \lesssim Y$ and $Y \lesssim X$, we write $X \approx Y$. Given a set of points, $P \subset [0, 1]^2$, let $\Pi_{\alpha, \beta}(P)$ denote the number of distinct triples of points, that determine a given pair of dot products. That is, for real numbers α and β ,

$$\Pi_{\alpha, \beta}(P) = \{(p, q, r) \in P \times P \times P : p \cdot q = \alpha \text{ and } p \cdot r = \beta\}.$$

We will typically restrict α and β to be positive nonzero, because zero dot products behave differently, as demonstrated by Proposition 2. Our first main result applies to any point set in the plane.

Theorem 1. *Given a set, P , of n points in \mathbb{R}^2 , and fixed $\alpha, \beta \neq 0$,*

$$|\Pi_{\alpha,\beta}(P)| \lesssim n^2.$$

In general, for a set of n points, $P \subset \mathbb{R}^2$, one cannot expect to get an upper bound better than Theorem 1, as shown in an explicit construction below, Proposition 1. In cases where the points are evenly-distributed, such as applications with sensor placement or code construction, we have tighter bounds on $\Pi_{\alpha,\beta}(P)$. Our second main result is for point sets with a minimum separation between points.

Theorem 2. *Let $P \subset [0, 1]^2$ be a set of n points that obeys the following separation condition*

$$\min\{|p - q| : p, q \in P, p \neq q\} \geq \epsilon.$$

For $\epsilon > 0$, and fixed $\alpha, \beta \neq 0$, we have

$$|\Pi_{\alpha,\beta}(P)| \lesssim n^{\frac{4}{3}} \epsilon^{-1} \log(\epsilon^{-1}).$$

Notice that if ϵ is chosen to be too small, Theorem 2 is outdone by Theorem 1. Similarly, if ϵ is close to 1, there cannot be many points in the unit square. Keeping this in mind, the range in which Theorem 2 is most useful is $n^{-\frac{2}{3}} < \epsilon \leq n^{-\frac{1}{2}}$. This range lines up with other results on a wide class of finite point sets called s -adaptable sets.

We now introduce the notion of s -adaptability. This should be viewed as a measure for how well-distributed the points are. This property has been used to study many types of geometric point configuration problems. See [6], [7], and [10], for example. Families of point sets which are s -adaptable can be used to transfer results between discrete point sets and sets with positive Hausdorff dimension. A large, finite point set $P \subset [0, 1]^2$, is said to be *s-adaptable* if the following two conditions hold:

$$\begin{aligned} (\text{energy}) \quad & \frac{1}{\binom{n}{2}} \sum_{\substack{p, q \in P \\ p \neq q}} |p - q|^{-s} \lesssim 1, \\ (\text{separation}) \quad & \min\{|p - q| : p, q \in P, p \neq q\} \geq n^{-\frac{1}{s}}. \end{aligned}$$

By setting $\epsilon = n^{-\frac{1}{s}}$, and appealing to the definition of s -adaptability given here, we get the following estimate as a corollary.

Theorem 3. *Let $P \subset [0, 1]^2$ be a set of n points that is s -adaptable. For $2 \geq s > \frac{3}{2}$, and fixed $\alpha, \beta > 0$,*

$$|\Pi_{\alpha,\beta}(P)| \lesssim n^{\frac{4}{3} + \frac{1}{s}} \log n.$$

First, we construct examples of point sets that illustrating the sharpness of Theorem 1 as well as an illustration of why we assume the restriction of $\alpha, \beta \neq 0$. Next, we prove Theorem 1 and Theorem 2 in Section 3. Section 4 contains the proofs of two technical lemmas, included for completion.

2. EXPLICIT CONSTRUCTIONS

2.1. Sharpness of Theorem 1.

Proposition 1. *Given a natural number n , and real numbers $0 < \alpha, \beta < 2$, there is a set, P , of n points in $[0, 1]^2$ for which*

$$|\Pi_{\alpha,\beta}(P)| \approx n^2.$$

Proof. Let p be the point with coordinates $(1, 1)$. Now, staying within the unit square, distribute $\lfloor \frac{n-1}{2} \rfloor$ points along the line $y = \alpha - x$, and distribute the remaining $\lceil \frac{n-1}{2} \rceil$ points along the line $y = \beta - x$. Clearly, there are $\gtrsim n^2$ pairs of points (q, r) , where q is chosen from the first line, and r is chosen from the second. Notice that p contributes a triple to $\Pi_{\alpha,\beta}(P)$ for each such pair, giving us

$$|\Pi_{\alpha,\beta}(P)| \approx n^2.$$

□

2.2. The special case $\alpha = \beta = 0$.

Proposition 2. *There exists a set, P , of n points in $[0, 1]^2$ for which*

$$|\Pi_{0,0}(P)| \approx n^3.$$

Proof. Arrange $\frac{n}{2}$ along the x -axis, and $\frac{n}{2}$ points along the y -axis. Now, for each of the $\frac{n}{2}$ points on the x -axis, there are $(\frac{n}{2})(\frac{n}{2})$ pairs of points on the y -axis. Notice that any point chosen from the x -axis will have dot product zero with each point from the pair chosen from the y -axis. Therefore, each of these $\frac{1}{8}n^3$ triples will contribute to $\Pi_{0,0}(P)$.

We can get just as many triples that contribute to $\Pi_{0,0}(P)$ by taking single points from the y -axis, and pairs of points from the x -axis. In total, we get

$$|\Pi_{0,0}(P)| = \frac{1}{8}n^3 + \frac{1}{8}n^3 \approx n^3.$$

□

3. PROOFS OF MAIN RESULTS

3.1. Proof of Theorem 1.

Proof. The basic idea is to estimate the number of triples, $(p, q, r) \in \Pi_{\alpha,\beta}(P)$, by considering pairs of points, $(q, r) \in P \times P$, and bounding the number of possible points, $p \in P$, which could contribute to $\Pi_{\alpha,\beta}(P)$.

Call any line through the origin a *radial line*. Let $p \in P$ have coordinates (p_x, p_y) . We define the α -line for a point, p , to be the set of points that have dot product α with the point p . Holding p fixed, this set will be a line of points q , with coordinates (q_x, q_y) satisfying the equation

$$(3.1) \quad p \cdot q = p_x q_x + p_y q_y = \alpha.$$

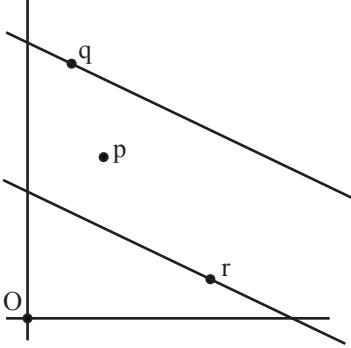


FIGURE 1: Here, q lies on the α -line of p , and r lies on the β -line of p .

We can define a β -line similarly. By solving (3.1) for q_y , we can see that the slope of the α -line of a point, p , will be equal to the slope of the β -line of the point p . Moreover, these lines will be perpendicular to the radial line through p . We let $\mathcal{L}_\alpha(p)$ and $\mathcal{L}_\beta(p)$ denote the set of all points of P incident to the α -line or β -line for the given point p .

Draw α -lines and β -lines for any given pair of $(q, r) \in P \times P$. We will refer to the set of point pairs with distinct α -lines and β -lines as A and pairs with a shared line as B . To be precise:

$$A = \{(q, r) \in P \times P : \mathcal{L}_\alpha(q) \neq \mathcal{L}_\beta(r) \text{ and } \mathcal{L}_\beta(q) \neq \mathcal{L}_\alpha(r)\},$$

and $B = (P \times P) \setminus A$.

We first consider triples in $\Pi_{\alpha,\beta}(P)$ of the form (p, q, r) where $(q, r) \in A$. Notice that the α -lines and β -lines of a pair of points, $(q, r) \in A$, can intersect at most four times, by definition of the set A . So for every pair of points in A , there are at most four possible locations for a point in P which would contribute a triple to $\Pi_{\alpha,\beta}(P)$. As $A \subset P \times P$, we see that $|A| \leq n^2$. From this, we see that pairs in A cannot add more than $4n^2$ triples to $\Pi_{\alpha,\beta}(P)$. //

We now turn our attention to the set B . Without loss of generality, suppose that the α -line of a point, $q \in P$, coincides with the β -line of a point, $r \in P$, then we appeal to the following lemma.

Lemma 3.1. *For any pair of points, $(q, r) \in B$, as defined above, the following hold:*

- (1) *Both q and r must lay along the same radial line.*
- (2) *The ratio of the distances from q and r to the origin must equal the ratio between β and α .*

Lemma 3.1 is proved in Section 4. With these conditions in tow, we see that the pairs of points in B are quite rare. Fix any radial line, L . Each point from $P \cap L$ can have its α -line overlap with at most one β -line. Similarly, each point from $P \cap L$ can have its β -line overlap with at most one α -line. So each point from $P \cap L$ can be in at most two pairs from B .

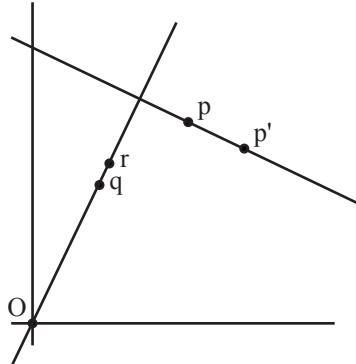


FIGURE 2: Here, the pair (q, r) lies on a radial line. The points p and p' lie on the α -line of q , which coincides with the β -line of r .

Any pair, $(q, r) \in B$, that lies on L will contribute as many triples of points, (p, q, r) , to $\Pi_{\alpha, \beta}(P)$ as there are points coincident to both the α -line of q and the β -line of r . As these families of shared lines are parallel for point pairs along L , each point, p , can be on at most one α -line, regardless of a possible overlap with a β -line. This means that each point $p \in P$ can be in at most one triple of the form $(p, q, r) \in \Pi_{\alpha, \beta}(P)$ with a pair of points, q and r , from L . The total number of triples contributed by pairs of points in L is therefore no more than n .

As there are no more than n points, there can be no more than n distinct radial lines to consider. Since each radial line can contribute no more than n triples to $\Pi_{\alpha, \beta}(P)$, the maximum contribution to $\Pi_{\alpha, \beta}(P)$ by pairs in B is no more than n^2 . \square

3.2. Proof of Theorem 2.

Proof. First, we define α -lines and β -lines as in the proof of Theorem 1. Again, we let $\mathcal{L}_\alpha(p)$ and $\mathcal{L}_\beta(p)$ denote the set of all points incident to the α -line or β -line for the given point p . Now, consider the α -lines and β -lines for each $p \in P$. Referring back to the definition of $\Pi_{\alpha, \beta}(P)$, we see that a triple of points, (p, q, r) , will be in $\Pi_{\alpha, \beta}(P)$ precisely when q lies on the α -line of p and r lies on the β -line of p . So the quantity we aim to estimate is

$$(3.2) \quad |\Pi_{\alpha, \beta}(P)| = \sum_{p \in P} |\mathcal{L}_\alpha(p)| |\mathcal{L}_\beta(p)|.$$

It follows that we want a bound on the number of times that a point, $p \in P$, is incident to an α -line or β -line. The following result will be proved in Section 4.

Lemma 3.2. *In the setting above, the number of point-line incidences, I , is*

$$I \lesssim n^{\frac{4}{3}}.$$

Recalling the definitions of $\mathcal{L}_\alpha(p)$ and $\mathcal{L}_\beta(p)$, we write I as:

$$I = \sum_{p \in P} (|\mathcal{L}_\alpha(p)| + |\mathcal{L}_\beta(p)|)$$

We now dyadically decompose P into two families of disjoint sets defined by the number of incidences they contribute.

$$P_j^\alpha := \{p \in P : 2^j \leq |\mathcal{L}_\alpha(p)| < 2^{j+1}\}$$

Where P_k^β is defined similarly. Now, if a point has roughly 2^j points from P on its α -line, and 2^k points from P on its β -line, then it will be in the intersection of P_j^α and P_k^β . Let us define these intersections as:

$$P_{j,k} := P_j^\alpha \cap P_k^\beta.$$

The intersection of any α -line or β -line with $[0, 1]^2$ can be no longer than $\sqrt{2}$. All of our points are contained in $[0, 1]^2$, so it will suffice to estimate the maximum number of points of P on any line segment of length $\leq \sqrt{2}$. Fix such a segment, ℓ . Recall the separation condition,

$$\min\{|p - q| : p, q \in P, p \neq q\} \geq \epsilon.$$

So each point on ℓ must have a vacant length of segment at least ϵ long in either direction. So if ℓ had points packed on it maximally, there would be no more than

$$\frac{\sqrt{2}}{\epsilon}.$$

We can see that for every point $p \in P$,

$$(3.3) \quad |\mathcal{L}_\alpha(p)| \lesssim \epsilon^{-1} \text{ and } |\mathcal{L}_\beta(p)| \lesssim \epsilon^{-1},$$

so we can be assured that $P_{j,k}$ is empty for j or k bigger than $\lceil \log_2(\epsilon^{-1}) \rceil$. This also tells us that for all relevant indices j and k in the sums to follow, we have:

$$(3.4) \quad 2^j, 2^k \lesssim \epsilon^{-1}$$

By combining the above:

$$\begin{aligned} n^{\frac{4}{3}} &\gtrsim I = \sum_{p \in P} (|\mathcal{L}_\alpha(p)| + |\mathcal{L}_\beta(p)|) \\ &\approx \sum_{j=0}^{\lceil \log_2(\epsilon^{-1}) \rceil} \sum_{k=0}^{\lceil \log_2(\epsilon^{-1}) \rceil} \left(\sum_{p \in P_{j,k}} (2^j + 2^k) \right) \\ &= \sum_{j=0}^{\lceil \log_2(\epsilon^{-1}) \rceil} \sum_{k=0}^{\lceil \log_2(\epsilon^{-1}) \rceil} (|P_{j,k}| (2^j + 2^k)). \end{aligned}$$

So, for any pair of indices, j and k , we have the following bound

$$(3.5) \quad |P_{j,k}| (2^j + 2^k) \lesssim n^{\frac{4}{3}}.$$

We now dyadically decompose the sum in (3.2) as we did with the sum estimating I .

$$\begin{aligned}
|\Pi_{\alpha,\beta}(P)| &= \sum_{p \in P} |\mathcal{L}_\alpha(p)||\mathcal{L}_\beta(p)| \\
(3.6) \quad &\approx \sum_{j=0}^{\lceil \log_2(\epsilon^{-1}) \rceil} \sum_{k=0}^{\lceil \log_2(\epsilon^{-1}) \rceil} (|P_{j,k}| (2^j) (2^k))
\end{aligned}$$

Let l and m be a pair of indices that give the largest summand in (3.6). Now we have:

$$\begin{aligned}
|\Pi_{\alpha,\beta}(P)| &\approx \sum_{j=0}^{\lceil \log_2(\epsilon^{-1}) \rceil} \sum_{k=0}^{\lceil \log_2(\epsilon^{-1}) \rceil} (|P_{j,k}| (2^j) (2^k)) \\
&\lesssim \sum_{j=0}^{\lceil \log_2(\epsilon^{-1}) \rceil} \sum_{k=0}^{\lceil \log_2(\epsilon^{-1}) \rceil} (|P_{l,m}| (2^l) (2^m)) \\
(3.7) \quad &\lesssim \lceil \log_2(\epsilon^{-1}) \rceil \lceil \log_2(\epsilon^{-1}) \rceil (|P_{l,m}| (2^l) (2^m))
\end{aligned}$$

Burying the constants from the logarithms in (3.7), we get:

$$(3.8) \quad |\Pi_{\alpha,\beta}(P)| \lesssim (|P_{l,m}| (2^l) (2^m)) \log \epsilon^{-1}.$$

Finally, by (3.8), adding in a term of 2^m , (3.5), and (3.4), we conclude

$$\begin{aligned}
|\Pi_{\alpha,\beta}(P)| &\lesssim (|P_{l,m}| (2^l) (2^m)) \log (\epsilon^{-1}) \\
&\lesssim (|P_{l,m}| (2^l + 2^m) (2^m)) \log (\epsilon^{-1}) \\
&\lesssim n^{\frac{4}{3}} (2^m) \log (\epsilon^{-1}) \\
&\lesssim n^{\frac{4}{3}} \epsilon^{-1} \log (\epsilon^{-1}),
\end{aligned}$$

as desired. \square

4. PROOFS OF LEMMAS

4.1. Proof of Lemma 3.1. To see (1), notice that any line that yields the prescribed dot products for a given point is perpendicular to that point's radial line. For the α -line of q and β -line of r to be coincidental, they must both be perpendicular to the same radial line, and thus, be generated from two points upon the same radial line.

For (2), assuming the α -line of q is coincidental to the β -line of r , we get the following two equations of lines in the plane:

$$\begin{aligned}
y &= \frac{\alpha}{q_2} - \frac{q_1}{q_2}x \\
y &= \frac{\beta}{r_2} - \frac{r_1}{r_2}x.
\end{aligned}$$

We know that q and r are on the same radial line, so there must exist a $\lambda > 0$ such that,

$$q_1 = \lambda r_1, \text{ and } q_2 = \lambda r_2.$$

We set the equations equal to one another and get:

$$\frac{\alpha}{q_2} - \frac{q_1}{q_2}x = \frac{\beta}{r_2} - \frac{r_1}{r_2}x.$$

By substituting for the coordinates of q ,

$$\frac{\lambda\alpha}{r_2} - \frac{\lambda r_1}{\lambda r_2}x = \frac{\beta}{r_2} - \frac{r_1}{r_2}x.$$

After simplifying, we see that

$$\lambda\alpha = \beta.$$

4.2. Proof of Lemma 3.2. There is a small technical obstruction to a direct application of the celebrated Szemerédi-Trotter point-line incidence theorem, which is that two types of line may be coincident. This turns out to not be a problem in our case. By appealing to the definitions of α -line and β -line, we can see that no two points can determine the same α -line or β -line. However, as we have seen above, it is possible for the α -line of a point to overlap the β -line of a point (a different point, unless $\alpha = \beta$).

We prove Lemma 3.2 using techniques of Székely, from [13]. Draw a graph (possibly a multigraph) with the n points as vertices, and the segments of the α -lines and β -lines connecting adjacent points as edges. Observe that no pair of points can have more than two edges connecting them, and this can only happen if the α -line of a point is coincident to the β -line of a point. If this does happen, we can replace the line segments by two curves, whose endpoints are the adjacent points in question, drawn in such a way so as not to cross any more or fewer edges than the initial line segment crossed.

Notice that the number of edges contributed by each line is equal to the number of point-line incidences on that line minus one (For example, if there are seven points on a given line, we would connect consecutive points with six edges.). So the number of edges is equal to the number of incidences minus the number of lines that contribute incidences. There are exactly $2n$ lines (n of each type) that may contribute incidences. Let I denote the number of point-line incidences, and e denote the number of edges in the graph. We have that

$$(4.1) \quad e \geq I - 2n.$$

We let $cr(G)$ denote the *crossing number* of G , which is the maximum number times that edges of G must cross one another at some point which is not a vertex, for any redrawing of G . We appeal to the crossing number lemma in [13].

Lemma 4.1. *Given a topological multi-graph, G , with v vertices, e edges, and a maximum edge multiplicity of m , if $e > 5mv$, then*

$$cr(G) \gtrsim \frac{e^3}{mv^2}.$$

For our setup, we have that $m \leq 2$, as there can be no more than two edges between any given pair of points. Now, either $e \leq 5mv$ or $e > 5mv$. In the first case, we recall (4.1) to see that

$$I \leq e + 2n \leq 5mv + 2n \leq 10n + 2n = 12n.$$

In the case that $e > 5mv$, we appeal to Lemma 4.1 and get

$$\frac{(I - 2n)^3}{2n^2} \leq \frac{e^3}{mv^2} \lesssim cr(G).$$

Notice that the crossing number of the graph can be no more than the number of times that the α -lines and β -lines crossed one another. Since there are $2n$ total lines, and each line could potentially cross almost all of the others, the total number of crossing lines is bounded above by $(2n)^2$. As this would correspond to a drawing of the graph, G , we are guaranteed that the crossing number of G is no more than $(2n)^2$. Comparing upper and lower bounds on $cr(G)$ yields

$$\frac{(I - 2n)^3}{2n^2} \lesssim n^2,$$

which gives us that

$$I \lesssim n^{\frac{4}{3}}.$$

In either case, the claimed estimate holds.

REFERENCES

- [1] J. Alvarez-Bermejo, J. A. Lopez-Ramos, J. Rosenthal, D. Schipani, *Managing key multicasting through orthogonal systems*, arXiv:1107.0586v2, (2015). [1](#)
- [2] P. Bahls, *Channel assignment on Cayley graphs*. J. Graph Theory, 67: 169–177, (2011). doi: 10.1002/jgt.20523 [1](#)
- [3] J. J. Benedetto and M. Fickus, *Finite normalized tight frames*, Adv. Comput. Math. 18, pp. 357–385 (2003). [1](#)
- [4] J. Garibaldi, A. Iosevich, and S. Senger, *Erdős distance problem*, AMS Student Library Series, 56, (2011). [1](#)
- [5] L. Guth and N. H. Katz, *On the Erdős distinct distance problem in the plane*, Annals of Math., Pages 155–190, Volume 181 (2015), Issue 1. [1](#)
- [6] A. Iosevich, H. Jorati, and I. Laba, *Geometric incidence theorems via Fourier analysis* Trans. Amer. Math. Soc. **361** (2009) 6595–6611. [2](#)
- [7] A. Iosevich, M. Rudnev, and I. Uriarte-Tuero, *Theory of dimension for large discrete sets and applications*, Math. Model. Nat. Phenom. 9 (2014), no. 5, 148–169. [2](#)
- [8] A. Iosevich, O. Roche-Newton, M. Rudnev, *On an application of the Guth-Katz theorem*, Mathematics Research Notices 18, no. 4, 1–7 (2014). [1](#)
- [9] A. Iosevich and S. Senger, *Orthogonal systems in vector spaces over finite fields*, Electronic J. of Combinatorics, Volume 15, December (2008). [1](#)
- [10] A. Iosevich and S. Senger, *Sharpness of Falconer's estimate in continuous and arithmetic settings, geometric incidence theorems and distribution of lattice points in convex domains*, arXiv:1006.1397, (2010). [2](#)
- [11] J. Spencer, E. Szemerédi, and W. T. Trotter. *Unit distances in the Euclidean plane* B. Bollobás, editor, “Graph Theory and Combinatorics”, pages 293–303, Academic Press, New York, NY, (1984). [1](#)
- [12] S. Steinerberger *A note on the number of different inner products generated by a finite set of vectors*, Discrete Mathematics, 310, (2010), pp. 1112–1117. [1](#)
- [13] L.A. Székely, *Crossing numbers and hard Erdős problems in discrete geometry*, Combin. Probab. Comput. **6** (1997), no. 3, 353–358. [8](#)

[14] T. V. Pham and L. A. Vinh *Orthogonal Systems in Vector Spaces over Finite Rings*, Electronic J. of Combinatorics, Volume 19, Issue 2 (2012). ¹

E-mail address: `poozzab@udel.edu`¹, `stevensenger@missouristate.edu`

¹Partially supported by a grant from the University of Delaware for summer undergraduate research.